

DTIC FILE COPY

2

AD-A219 008

**DISCOVERING PROBLEM SOLVING
STRATEGIES: WHAT HUMANS DO
AND MACHINES DON'T (YET)**

Technical Report AIP - 74

Kurt VanLehn

Department of Psychology
Carnegie Mellon University
Pittsburgh, PA 15213

**The Artificial Intelligence
and Psychology Project**

Departments of
Computer Science and Psychology
Carnegie Mellon University

Learning Research and Development Center
University of Pittsburgh

DTIC
ELECTE
MAR 14 1990
S B D

Approved for public release; distribution unlimited.

90 03 12 071

2

DISCOVERING PROBLEM SOLVING STRATEGIES: WHAT HUMANS DO AND MACHINES DON'T (YET)

Technical Report AIP - 74

Kurt VanLehn

Department of Psychology
Carnegie Mellon University
Pittsburgh, PA 15213

21 July 1989

DTIC
ELECTE
MAR 14 1990
S B D

This research was supported by the Computer Sciences Division, Office of Naval Research, under contract number N00014-86-K-0678. Reproduction in whole or part is permitted for any purpose of the United States Government. Approved for public release; distribution unlimited.

Unclassified

SECURITY CLASSIFICATION OF THIS PAGE

REPORT DOCUMENTATION PAGE

| | | | |
|---|--|--|----------------------------------|
| 1a. REPORT SECURITY CLASSIFICATION Unclassified | | 1b. RESTRICTIVE MARKINGS | |
| 2a. SECURITY CLASSIFICATION AUTHORITY | | 3. DISTRIBUTION / AVAILABILITY OF REPORT Approved for public release; Distribution unlimited | |
| 2b. DECLASSIFICATION / DOWNGRADING SCHEDULE | | | |
| 4. PERFORMING ORGANIZATION REPORT NUMBER(S) AIP - 74 | | 5. MONITORING ORGANIZATION REPORT NUMBER(S) | |
| 6a. NAME OF PERFORMING ORGANIZATION Carnegie-Mellon University | 6b. OFFICE SYMBOL (if applicable) | 7a. NAME OF MONITORING ORGANIZATION Computer Sciences Division Office of Naval Research | |
| 6c. ADDRESS (City, State, and ZIP Code) Department of Psychology Pittsburgh, Pennsylvania 15213 | | 7b. ADDRESS (City, State, and ZIP Code) 800 N. Quincy Street Arlington, Virginia 22217-5000 | |
| 8a. NAME OF FUNDING / SPONSORING ORGANIZATION Same as Monitoring Organization | 8b. OFFICE SYMBOL (if applicable) | 9. PROCUREMENT INSTRUMENT IDENTIFICATION NUMBER N00014-86-K-0678 | |
| 8c. ADDRESS (City, State, and ZIP Code) | | 10. SOURCE OF FUNDING NUMBERS p4000ub201/7-4-86 | |
| | | PROGRAM ELEMENT NO N/A | PROJECT NO N/A |
| | | TASK NO. N/A | WORK UNIT ACCESSION NO N/A |
| 11. TITLE (Include Security Classification) Discovering problem solving strategies: What humans do and machines don't (yet) | | | |
| 12. PERSONAL AUTHOR(S) VanLehn, Kurt | | | |
| 13a. TYPE OF REPORT Technical | 13b. TIME COVERED FROM 86Sept15 TO 91Sept14 | 14. DATE OF REPORT (Year, Month, Day) 1989 July 21 | 15. PAGE COUNT 14 |
| 16. SUPPLEMENTARY NOTATION Submitted to Sixth International Workshop on Machine Learning, for either the session on "Learning Plan Knowledge" or "Combining Empirical and Explanation-based learning" | | | |
| COSATI CODES | | 18. SUBJECT TERMS (Continue on reverse if necessary and identify by block number) | |
| FIELD | GROUP | SUB-GROUP | |
| | | Strategy discovery, skill acquisition, machine learning, cognitive science, impasses-driven learning, Artificial Intelligence, Problem-Solving, Human Factor | |
| 19. ABSTRACT (Continue on reverse if necessary and identify by block number) SEE REVERSE SIDE ENGINEERING (JG) | | | |
| 20. DISTRIBUTION / AVAILABILITY OF ABSTRACT <input type="checkbox"/> UNCLASSIFIED/UNLIMITED <input checked="" type="checkbox"/> SAME AS RPT <input type="checkbox"/> DTIC USERS | | 21. ABSTRACT SECURITY CLASSIFICATION | |
| 22a. NAME OF RESPONSIBLE INDIVIDUAL Dr. Alan L. Meyrowitz | | 22b. TELEPHONE (Include Area Code) (202) 696-4302 | 22c. OFFICE SYMBOL N00014 |

DD FORM 1473, 84 MAR

83 APR edition may be used until exhausted.

All other editions are obsolete.

SECURITY CLASSIFICATION OF THIS PAGE

Unclassified

People can discover new problem solving strategies on their own, without help from a teacher, text or other source. Many machine learning programs exist that discover strategies under similar conditions. Do we now have a sufficient set of computational models for understanding human strategy discoveries? This paper presents an unusually detailed analysis of a human problem solving protocol that uncovers 10 cases of strategies being discovered. It is argued that most cases are adequately modeled by existing machine learning techniques, and several are not, which suggests some interesting research problems for machine learning.

The claims are backed by a line-by-line simulation of the protocol using the Teton system (VanLehn & Ball, 1977; VanLehn, Ball & Kowalski, 1988), a descendent of Sierra (VanLehn, 1987; VanLehn, 1983). Some of Teton's strategy discoveries are provided by the user, as the technology for mechanizing them is not yet understood. This paper will not present Teton or its simulation of the protocol, since that information is available elsewhere (VanLehn, 1989). Instead, it will present the gist of the analysis, and point out its implications for machine learning.

The paper has five parts. After a brief discussion of the methods of the analysis and the protocol, the protocol analysis is presented in enough detail to allow evaluation of the accuracy of the empirical claims. A subsequent section classifies the cases of strategy discovery found in the data are classified according to standard machine learning concepts. The last section indicates which types of learning exhibited by the subject have not yet been exhibited by machine learning systems. This leads to the view that strategy acquisition by a competent human is like scientific theory formation, with the attendant tasks of experiment design and interpretation, noticing of serendipitous events, and even Eurisko-like hypothesis generation (Lenat & Brown, 1984). Although current machine learning models of strategy acquisition seem pale by comparison, there seems to be nothing stopping us from building machine learning systems with human-level capabilities for strategy discovery.

← cont next pt
 for keywords



| | |
|--------------------|--|
| Accession For | |
| NTIS GRA&I | <input checked="checked" type="checkbox"/> |
| DTIC TAB | <input type="checkbox"/> |
| Unannounced | <input type="checkbox"/> |
| Justification | |
| By | |
| Distribution/ | |
| Availability Codes | |
| Dist | Avail and/or Special |
| A-1 | |

People can discover new problem solving strategies on their own, without help from a teacher, text or other source. Many machine learning programs exist that discover strategies under similar conditions. Do we now have a sufficient set of computational models for understanding human strategy discoveries? This paper presents an unusually detailed analysis of a human problem solving protocol that uncovers 10 cases of strategies being discovered. It is argued that most cases are adequately modeled by existing machine learning techniques, and several are not, which suggests some interesting research problems for machine learning.

The claims are backed by a line-by-line simulation of the protocol using the Teton system (VanLehn & Ball, 19??; VanLehn, Ball & Kowalski, 1988), a descendent of Sierra (VanLehn, 1987; VanLehn, 1983). Some of Teton's strategy discoveries are provided by the user, as the technology for mechanizing them is not yet understood. This paper will not present Teton or its simulation of the protocol, since that information is available elsewhere (VanLehn, 1989). Instead, it will present the gist of the analysis, and point out its implications for machine learning.

The paper has five parts. After a brief discussion of the methods of the analysis and the protocol, the protocol analysis is presented in enough detail to allow evaluation of the accuracy of the empirical claims. A subsequent section classifies the cases of strategy discovery found in the data are classified according to standard machine learning concepts. The last section indicates which types of learning exhibited by the subject have not yet been exhibited by machine learning systems. This leads to the view that strategy acquisition by a competent human is like scientific theory formation, with the attendant tasks of experiment design and interpretation, noticing of serendipitous events, and even Eurisko-like hypothesis generation (Lenat & Brown, 1984). Although current machine learning models of strategy acquisition seem pale by comparison, there seems to be nothing stopping us from building machine learning systems with human-level capabilities for strategy discovery.

1. Methodological preliminaries

In protocol data, subjects rarely announce that they have discovered of a new rule or concept, which has led to the impression that human learning is mostly a gradual, automatic compilation of productions (Anderson, 1983) or chunks (Newell, 19??) from knowledge obtained by reading or hearing instructions. Although such mechanisms surely exist in the human cognitive architecture, and may in fact be the only underlying mechanisms for storing new items in long term memory, there are a number of phenomena that they do not explain. For instance, some people learn much more quickly than others -- how can this occur if they all have the same cognitive architecture? A common view is that better learners have strategies that they habitually use to detect deficiencies in their knowledge and seek remedies. For instance, Chi et al. (in press) have shown that good students studying physics examples try to explain the example to themselves whereas poor students merely rehearse or paraphrase the example. This higher level view of learning, which emphasizes the more-or-less conscious application of learning strategies, is the major pretheoretical bias in the present investigation.

The protocol analysis is founded on the assumption that learning events (i.e., the application of a learning strategy) are much more common in human behavior than has previously been supposed, but subjects rarely mention them. Instead, they show up mostly as pauses in the protocol data. Thus, it takes detailed, line-by-line analysis and computer simulation in order to detect the abrupt, but subtle shifts in strategy that characterize the acquisition of a new strategic rule. Line-by-line analysis is extremely tedious and rarely used. Although Newell and Simon (1972) used this method to analyze problem solving (they specifically ignored learning), it has not appeared in the cognitive science literature since then. In part, this research demonstrates the feasibility of using line-by-line analysis for locating learning events. This gives the field of machine learning a tool for determining which of its many techniques for learning correspond to human learning, and more importantly, for uncovering facts of human learning that cannot yet be duplicated by machines. Although machines can learn more than they ever did, humans are still better learners than machines (or at least, the good human learners are).

A simple task domain, the Tower of Hanoi was chosen for two reasons. Although the long-term objective is to study events in the learning of physics (VanLehn, 19??a), there was enough uncertainty

about the feasibility of using line-by-line simulation to uncover learning events that a simpler task domain was chosen for initial investigation. Secondly, Anzai and Simon (1979) published a Tower of Hanoi protocol that exhibits significant amounts of learning and has unusually clear verbal statements by the subject. Their analysis of the protocol summarized the main strategies of the subject, but it did not locate the exact places in the protocol where strategic knowledge was acquired. This data is ideal for testing the feasibility of the analysis method, for if the method fails on these data, there is no hope for it on the physics protocols. Fortunately, the method did not fail, and we now have a preliminary list of learning events.

2. The protocol

In order to establish a context for discussing the learning events, the overall structure of the protocol needs to be presented. Throughout, the original line numberings and nomenclature of Anzai and Simon (1979) will be used so that readers may refer back to the published protocol for details. The pegs of the puzzle are labeled A, B and C, and the disks are numbered according to their size, with 1 being the smallest disk. The initial state of the puzzle has disks 1 through 5 on peg A. The goal is to get them all on peg C, subject to the constraints that a larger disk may never be placed on a smaller disk and only one disk may be moved at a time.

The protocol, which lasts 90 minutes, is divided into four major episodes. During episode 1, the subject starts solving the puzzle but gets the initial move wrong and eventually gives up. During episode 2, the subject deliberately selects the other legally possible initial move, and succeeds in solving the puzzle along the optimal solution path. However, she struggles at many points to find the correct move, apparently because she is looking ahead several moves in her mind's eye in order to evaluate moves before making them. She is apparently not satisfied with this strategy, so during episode 3, she embarks on a "experiment," wherein she successively solves increasingly larger versions of the puzzle. She starts by solving the trivial puzzle that has just one disk on peg A. Then she solves the puzzle whose initial state has two disks on peg A, and so on. Most of her learning occurs during this episode, and she emerges with a clear strategy based on recursive subgoalings. She seems quite satisfied with this, and only solves the puzzle again, in episode 4, because the experimenter asks her to. Nonetheless, a subtle strategic shift occurs during episode 4.

3. The learning events

At the end of the protocol, the subject has, according to our analysis, the rules shown in table 3-1. If the rule was acquired during the protocol, the line number of the rule's learning event appears in brackets after the rule. The first six rules form a "partial" strategy that develops early and is used throughout the rest of the protocol. This strategy uniquely determines 26 of the 31 moves along the subject's path. At each of the other 5 moves (moves 1, 5, 9, 17 and 25), the strategy is ambiguous, and allows the subject a choice of two legal disk movements. For easy future reference, let us call moves 1, 5, 9, 17 and 25 the *major moves* of the solution path, and use *minor move* to refer to the other 26 moves. The first six rules will be called the *minor move strategy*. All the other rules are used for determining major moves. The major moves are where we see the largest number of learning events as the subject invents increasingly more effective strategies for making the major moves.

The next few sections discuss the rules of table 3-1 and characterize their learning events in some detail. The casual reader may wish to skip ahead to the section "classification of the learning events."

3.1. The initial rules

The subject seems to acquire rules 1 through 4 by reading the instructions to the puzzle, because these rules are used without comment at the first possible occasion where they can be applied. Notice that rule 4 is common sense: if something is in the way, move it out of the way. However, it contains the seeds of strategies that the subject eventually learns. Rule 4, however, applies only when a single disk is blocking the move. The later strategies apply no matter how many disks are blocking the move.

-
1. Achieve the top level goals of the puzzle in the following order: get disk 5 to C, get disk 4 to C, get disk 3 to C, get disk 2 to C and get disk 1 to C.
 2. Do not move the same disk on consecutive moves.
 3. If there is a choice of where to put disk 1, and disk 2 is exposed, then put disk 1 on top of disk 2, thus creating a small pyramid.
 4. If the goal is to move a given disk from a given peg to another given peg, and there is exactly one disk blocking the move, then get that blocking disk to the peg that is not involved in the move.
 5. Before working on achieving any of the top level goals, get disk 4 to peg B. [12]
 6. If the goal is to move a given disk from a given peg to another given peg, and the two-high pyramid is blocking the move, then get disk 1 to one of the two pegs involved in the move (thus allowing disk 2 to move out of the way of the move). [30-34]
 7. If the goal is to move disk 2 from peg A to peg C, and disk 1 is on peg A, then move disk 1 to the peg that is not involved in the move. [78]
 8. If the goal is to move disk N from peg A to peg C, and disk N-1 is on peg A, then get disk N-1 to the peg that is not involved in the move. [82]
 9. If the goal is to move disk N from peg A to a given peg, and disk N-1 is on peg A, then get disk N-1 to the peg that is not involved in the move. [84]
 10. If the goal is to move disk N from a given peg S to a given peg T, and disk N-1 is on S, then get disk N-1 to the peg that is not involved in the move. [99]
 11. If the goal is to move a given disk from a given peg to another given peg, and disk D is the largest disk blocking the move, then get D to the peg that is not involved in the move. [121]
 12. If the goal is to move a given pyramid from a given peg to another given peg, and pyramid P is the largest pyramid blocking the move, then get P on the peg that is not involved in the move. [179]

Table 3-1: Rules used by the subject during the protocol

3.2. Learning rule 5

Rule 5 seems to be acquired at lines 11-12. At line 11, the subject has gotten the puzzle into the state [45,123,_]. (This notation means that disks 4 and 5 are on peg A, disks 1, 2 and 3 are on peg B, and peg C is empty.) In lines 11-13, she says "So then, 4 will go from A to C. And then..., um..., oh..., um..., I should have placed 5 on C." During the long pause, the subject seems to realize that placing disk 5 on peg C means that disk 4 has to go to peg B first. She does not seem to form a general recursive subgoal rule (that comes later), but instead changes her representation of the goals for the puzzle by prefixing the goal of getting disk 4 to peg B. This explains why she is operating with that goal at lines 30-34. In episodes 3 and 4, at the lines parallel to this move (119-124 and 173-174, respectively), where she is using her new recursive subgoal strategies, she starts with the goal of getting 4 to B, rather than the goal of getting 5 to C. This makes it seem quite likely that she has merely "rotely memorized" the 4-to-B goal rather than forming a general, recursive rule.

The acquisition of the 4-to-B goal seems to be triggered by an impasse.¹ Just after moving from state [45,123,_] to [5,123,4], the subject appears to get stuck. One explanation of this impasse is that she focuses on the goal of putting disk 5 on peg C in the state [5,123,4] where only disk 4 blocks the move. This triggers her common sense rule about moving an object out of the way (rule 4), so she formulates the subgoal of moving disk 4 to peg B. This subgoal cannot be immediately achieved, because peg B is occupied by smaller disks than 4. Thus, she is at an impasse.²

The subject's comment, "I should have placed 5 on C," sheds no light on the type of reasoning she is during this learning event. Since this is her first experience with the puzzle, and similarity-based learning works best when the learner has multiple events to generalize over, it is slightly more plausible that she uses explanation-based learning here. For instance, she could have reasoned:

- Because getting disk 4 out of the way on B will always be a subgoal of moving disk 5 from A to C (by rule 4),
- and because moving disk 5 to C is always the first top level goal to be achieved in the five-disk puzzle (by rule 1),
- it follows that moving disk 4 to B is a prerequisite to achievement of any of the top level goals of the five-disk puzzle.

Whether or not the subject makes this deduction is unclear from the protocol. However, it is fairly clear that she does it in response to an impasse.

3.3. Learning rule 6

Rule 6 is "If the goal is to move a disk from one peg to another, and the two-high pyramid is blocking the move, then move disk 1 to one of the pegs involved in the move so that disk 2 can move out of the way." The subject seems to acquire this rule at lines 30-34. This is the first occasion in the whole protocol where a two-high pyramid is blocking a desired move. As we shall see in a moment, she does not handle the problem smoothly, indicating that the rule had not been acquired prior to lines 30-34. After line 34, there are many occasions for rule 6 to apply, and it seems to be involved in most of them. After line 34, there are 3 occasions in episode 2 where the rule could apply. On 2 of them (lines 43 and 58), the subject utters only the perfunctory X-from-Y-to-Z comment. On the last one (line 63), the subject says "This time, if I think of 3 on C, that will be good, so 1 will go from A to C." Thus, on all three occasions, the subject seems to execute rule 6 with no difficulty. There are 9 more occasions later in the protocol where rule 6 could apply. However, rule 6 competes with the subgoaling strategy for command of these moves. On four of the moves (lines 92, 136, 148 and 184), the subject utters the usual X-from-Y-to-Z comment, which can be interpreted as the firing of rule 6. On the remaining five moves (lines 79-85, 96-101, 153-155, 204-205 and 210-212), the subject makes the usual comments appropriate for the recursive subgoaling strategies (i.e., lines 153-155: "3 naturally has to go here, so, for that, 2 has to go to B. So 1 will go from A to C."). In summary, the evidence for rule 6 being learned at lines 30-34 is that it could not have been learned earlier and it seems to be used regularly thereafter.

Let us now see what the subject says during this critical segment in order to infer how she learned rule 6. (In this and subsequent excerpts from the protocol, the state of the puzzle is shown in the first column.) The segment where rule 6 is learned is:

- | | | |
|-----------|-----|--|
| [45,12,3] | 30. | And so I'll place 1 from B... to C. |
| | 31. | Oh yeah! I have to place it on C |
| | 32. | Disk 2... no, not 2, but I placed 1 from B to C... Right? |
| | 33. | Oh, I'll place 1 from B to A. (Experimenter: Go Ahead.) |
| [145,2,3] | 34. | Because... I want 4 on B, and if I had placed 1 on C from B, it wouldn't have been able to move. |

Apparently, the subject has the goal of moving disk 4 from A to B (line 34), just as rule 5 predicts. However, that move is blocked by the two-high pyramid on peg B. She starts by applying rule four, which moves a blocking disk out of the way (line 30). She imagines that she has moved disk 1 from B to C but does not actually move it.³ Line 31 seems to be a repeat of the inference that a blocking disk must be moved out of the way. Repeating an inference is common in protocol data (see Newell and Simon, 1972). At line 32, she continues looking ahead in her mind's eye, visualizing where she will put disk 2. Disk 2 can only move to peg A, which means that disk 4 is now blocked by two disks. She seems to reach an impasse at this point. Her resolution of the impasse comes in line 33, where she backs up to her earlier decision about where to move disk 1. She reverses her decision, and moves disk 1 to peg A. In line 34, she double-checks her reasoning. Line 34 seems to be where she forms rule 6.

If this interpretation of the protocol is correct, then the acquisition of rule 6 is triggered by an

impasse. The impasse is caused by taking two moves (or rather, visualizing them) in order to get disk 4 free to move to B, then discovering that disk 4 is still blocked.

The type of reasoning going on after the impasse is clear enough, but it is unclear how to classify it. The subject first applies the repair strategy of backing up to a previous choice point and taking the other choice (the back-up repair of Repair Theory -- Brown and VanLehn, 1980; VanLehn, 1983; VanLehn, in press). This results in moving 1 to A. After that, she re-examines her reasoning in line 34. However, because neither her repair strategy nor her re-examination of the repair are based on any domain-specific rules, this type of reasoning does not qualify as explanation-based learning. On the other hand, it is clearly not similarity-based learning, which would involve looking for similarities and dissimilarities across several situations. So neither of the standard classifications for learning events seems to fit the subject's reasoning.

On the other hand, the subject's reasoning fits beautifully with a type of learning conjectured by Brown and VanLehn (1980) and modeled by Sierra (VanLehn, 1983; VanLehn, 19??b). They conjectured that subjects would sometimes acquire stable buggy procedures by applying a repair strategy and then storing the results in memory in such a way that the next time this impasse occurred, exactly the same actions would be taken. Brown and VanLehn called this type of learning *patching*, because that is a term computer programmers use for fixing a program in an unprincipled, superficial way. This learning event seems to be a clear case of patching. The subject executes a repair strategy, then looks over the results. Apparently, this suffices to form a stable rule (rule 6).⁴

3.4. Learning disk subgoalting: Rules 7 through 11

In its most general form, the disk subgoalting strategy is: if the goal is to move a disk from a given peg to another peg, and disk D is the largest disk blocking the move, then get disk D to the peg that is not involved in the move. This strategy is recursive, and a specialization of operator subgoalting, a venerable weak method. As an illustration of the strategy's operation, at lines 108-114 the subject says, when planning the initial move of the five disk puzzle, "5 will have to go to C, right? So 4 will be at B. 3 will be at C. 2 will be at B. So 1 will go from A to C." and then makes her initial move, which is to place disk 1 on peg C. This strategy is called *disk subgoalting* in order to distinguish it from a similar recursive strategy, discussed later, that uses pyramids instead of disks as the objects being reasoned about.

There is no evidence of the disk subgoalting strategy in episode 2, but by the time the subject starts on the five-disk puzzle in episode 3, there is ample evidence of the strategy (the lines just quoted are the first move of five-disk puzzle's solution in episode 3). The strategy seems to develop during the first part of episode 3, while the subject is engaged in her experiment of trying to solve towers of increasing height.

The acquisition of the strategy takes five learning events. The first one generates an initial version of the rule (rule 7), and the subsequent learning events generalize it, eventually producing a fully general strategy (rule 11). The subject begins by setting up the trivial one-disk puzzle and solving it, saying:

[1,_,_] 76. First, if I think of it as only one disk, 1 could go from A to C, right?

Nothing seems to be learned from this experiment, so she sets up the two-disk puzzle and solves it, saying:

77a. But, if you think of it as two disks,
[12,_,_] 77b. this will certainly go as 1 from A to B
[2,1,_] 77c. and 2 from A to C,
[_,1,2] 77d. then 1 from B to C.
[_,_,12] 78a. That...
78b. that anyway, 2 will have to go to the bottom of C,
78c. naturally I thought of 1 going to B.

At line 77a, she is setting up the puzzle. In 77a, 77b and 77c, she solves the puzzle smoothly, apparently by applying the minor move strategy. However, she does not proceed immediately to solving the three disk puzzle, but pauses at line 78a and reflects on what she has just done.

Apparently, the subject is not just solving these puzzles for the sake of practice or amusement. Rather, she seems to have a plan that consists of solving each version of the puzzle then reflecting on the solution in order to find a better solution strategy (or just understand the puzzle better). Moreover, the subject has deliberately varied the number of disks in the puzzle (she could, for instance, have varied the starting peg instead). This suggests that she is looking for a rule that is independent of the number of disks in the puzzle and of the identity of particular disks. For instance, her rule for moving the two-high pyramid out of the way (rule 6) is exactly the kind of rule she wants to avoid, for it mentions particular disks. In a moment, evidence will be presented that she is deliberately ignoring rule 6 throughout this section of the protocol, probably because she wants to find a more powerful rule.

Her reflection at line 78a results in the comments at 78b and 78c. The word "naturally" indicates that she sees a connection between moving 1 to B and the goal of getting 2 on C. Although it could be that this connection comes from rule 4, the subsequent protocol evidence indicates that this is the beginning of the acquisition of the disk subgoal strategy. In particular, she seems to form the following rule (rule 7): "If the goal is to move disk 2 from peg A to peg C, and disk 1 is on peg A, then move disk 1 to the peg that is not involved in the move." This is a very specific rule, which she later generalizes. Indeed, it is so specific that to call it a rule is a little presumptuous. Perhaps it should be called a "noteworthy experience" or a "situated rule."

In lines 79 through 85, the subject seems to make two attacks on the three-disk puzzle:

- [123,_,_] 79. So if there were three... yes, yes, now it gets difficult.
 90. Yes, it's not that easy...
 81. ...this time, 1 will...
 82. Oh, yes, 3 will have to go to C first.
 83. For that, 2 will have to go to B.
 84. For that, um... 1 will go to C.
 85. so, 1 will go from A to C.

Her first attack lasts from lines 79 to 81. She is probably working on the goal of moving the largest disk, disk 3, to peg C, as that is the goal selected by rule 1. However, she seems to suffer an impasse. This indicates that she did not learn a fully general version of the disk subgoal strategy back at line 78, since that would successfully reduce her goal and thus avoid the impasse. It also indicates that she is not using rule 6, because that rule readily applies to the state [123,_,_] and thus avoids the impasse. This is consistent with her supposed policy of ignoring rule 6 in order to find a more powerful rule.

How she handles the impasse is somewhat unclear. Since she mentions disk 1 in line 81, it may be that she is trying to look ahead in her mind's eye in order to determine which of the two legal moves of disk 1 leads to a goal. Since the solution path is 7 moves long for this puzzle, that attempt fails.

Regardless of the nature of her first attempt at resolving the impasse, her second attempt is to generalize the "situated rule" (or noteworthy experience, or whatever), and that attempt succeeds. At line 82, she says "Oh, yes," and proceeds to produce a classic disk goal-subgoal utterance. This suggests that she has generalized the situated rule of line 81 by substituting variables for the disk names, yielding rule 8: "If the goal is to move disk N from peg A to peg C, and disk N-1 is on peg A, then get disk N-1 to the peg that is not involved in the move." This rule applies to the state [123,_,_], whereas rule 7 does not. In short, lines 79-82 constitute a case of impasse-driven generalization of an existing rule.

At line 84, the subject pauses. This is consistent with the assumption that her rule was generalized by replacing the disk names with variables and leaving the peg names alone. Since the goal at line 84 is to move disk 2 from peg A to peg B, and rule 8 specifies "If the goal is to move disk N from peg A to peg C...", the rule does not immediately apply. This suggests that the pause at line 84 is due to an impasse, followed by another slight generalization of the rule. The rule produced by this learning event, rule 9, is: "If the goal is to move disk N from peg A to X, and disk N-1 is on peg A, then get disk N-1 to the peg that is not involved in the planned move."⁵

At this point, the subject seems to have acquired an initial version of the disk subgoal strategy. The strategy is highly specific in that it only works when the goal disk and the blocking disk are both on

peg A. The initial version of the strategy was acquired at line 78 via some kind of reflection. Twice it was generalized in response to impasses. This rule is sufficiently general to allow impasse-free problem solving of the three-disk puzzle and part of the four-disk puzzle. The rule fails to apply in the following state:

- | | | |
|-----------|------|---|
| [_,123,4] | 96. | And then, again this will go from A... I will... |
| | 97. | Wrong..., this is the problem and... |
| | 98. | 1 will go from B to C... |
| | 99. | For that, um... this time 3 from B, um... has to go on C, so... |
| | 100. | For that, 2 has to go to A. |
| | 101. | For that, 1 has to go back to C, of course. |

It is not clear why the subject mentions peg A in line 96. Perhaps she makes a mistake in applying some rule and catches herself at line 97. In the middle of line 97, she seems to start over. The lack of subgoaling commentary indicates a use of rule 6 in the derivation of the correct move, announced at line 98. For some reason, she decides to redo that derivation, starting at line 99, using her new subgoaling strategy. This is consistent with the proposal that the subgoaling strategy is invoked when the minor move strategy is failing, or at least when it seems to be failing, as in this case. It could also be that she is double-checking the old strategy against the new one. Any way, for whatever reason, it is clear that she begins to apply the new strategy at line 99.

However, the subgoaling rule will not immediately apply, because it is too specific. When we last saw it, the rule was: "If the goal is to move disk N from peg A to X, and disk N-1 is on peg A, then move disk N-1 to the peg that is not involved in the move." In this case, all the relevant disks are on peg B, so the rule does not apply. An impasse-driven generalization of the rule seems to be the source of the pauses in line 99. The new rule, rule 10, seems to be, "If the goal is to move disk N from peg Y to peg X, and disk N-1 is on peg Y, then move disk N-1 to the peg that is not involved in the move." This generalized rule is sufficient to produce the subgoaling seen in lines 99-101.

The subject finishes the four-disk puzzle and begins the five-disk puzzle uneventfully using a combination of the minor move strategy and rule 10. However, overspecificity of the rule causes an impasse during the following segment:

- | | | |
|--------------|------|--|
| [2345,_,1] | 116. | And then, 2 will go from A to B. |
| [345,2,1] | 117. | 1 will go back from B to C.[sic] |
| [345,12,_]] | 118. | 3 will go from A to C. |
| [45,12,3] | 119. | For that, um..., this time, again..., as this time 4 will have to go to B... |
| | 120. | Let's move back 1 from B to A.... |
| | 121. | If 4 has to go from A to B, it means... |
| | 122. | 2 will have to go to 3. |
| | 123. | Because 1 will.... |
| | 124. | So, 1 will go back from B to A. |

During lines 119-124, the subject waffles about using rule 6 of the minor move strategy. She seems to apply it at line 120, because she goes directly from the goal 4-to-B to the move of 1 to A without mention an intermediate subgoal. However, for some unknown reason (is she deliberately exercising her new strategy?), she begins at line 121 to rederive the move using rule 10. However, the rule is too specific to apply to the current situation. Rule 10 is "If the goal is to move disk N from peg S to peg T, and disk N-1 is on peg S, then move the blocking disk to the peg that is not involved in the move." This rule fails to apply for two reasons. First, the disk to be moved is disk 4 (i.e., N=4), but the largest blocking disk is only of size 2, not 3, as required by the rule. Second, the rule applies only when the blocking disks are on the same peg as the disk to be moved, but in the current situation, the blocking disks are on the peg B, which is the destination of the 4-to-B move.

The pause at the end of line 121 seems to be an impasse where the rule is generalized to become rule 11: "If the goal is to move a disk from one peg to another, and there are some disks blocking the move, then move the largest blocking disk to the peg that is not involved in the move." This version is just general enough to allow it to apply to the current puzzle state. Moreover, rule 11 is sufficiently general to apply to all puzzle states. The subject has at last acquired the disk subgoaling strategy. She finishes the

rest of the five-disk puzzle without any major impasses.

Summary

By line 164, the subject seems to have acquired a correct, smoothly functioning disk subgoal strategy. She began the acquisition of it way back at line 78. That initial version of the rule seems to have been very specific. This led to a succession of impasses (at lines 79, 84, 99, 121 and possibly 140), each of which caused a slight generalization of the rule.

The first learning event, at line 78, was not triggered by an impasse. Rather, the subject seems to have adopted a plan of solving successive versions of the puzzle and deliberately examining their solutions after each is made. The learning event at line 78 seems to be a case of deliberate reflection on the solution of the 2-high puzzle. The subject's reasoning during this learning event is a case of deduction. She says, "that anyway, 2 will have to go to the bottom of C, naturally I thought of 1 going to B." Her first clause is an application of the general idea that when one wants to build a stack of things, the one on the bottom must be placed first. Her second clause is an application of rule 4, which indicates how to get a single disk out of the way of a desired move. Thus, the learning event is a case of explanation-based learning triggered by a deliberate goal of reflecting on a solution.

The other learning events are all quite similar. They are triggered by impasses, and they produced slightly more general versions of the rule. The reasoning here does not seem to be deductive, because deduction could produce a fully general rule immediately.⁶ The gradual increase in specificity, where each step removes or abstracts the fewest features necessary in order to get the rule to match the new situation, is characteristic of induction. Indeed, exactly this policy of conservative generalization has been observed many times (e.g., Hunt, Marin & Stone, 1966; Smith & Medin, 1981; VanLehn, in press). So these learning events seem to be clear cases of similarity-based learning.

3.5. Learning the pyramid subgoal strategy

The pyramid subgoal strategy is similar to the disk subgoal strategy, except that the subject thinks in terms of pyramids instead of disks. A rule for pyramid subgoal can be derived from rule 11 for disk subgoal by merely substituting "pyramid" for "disk," yielding rule 12: "If the goal is to move a pyramid from a peg to another peg, and there are some pyramids blocking the move, then move the largest pyramid to the peg that is not involved in the move."

The pyramid subgoal strategy seems to be acquired somewhere in episode 4. At the first major move of episode 4, the subject's subgoals are disks: "Because, 5, at the end, will go to C, so, so, 4 will go to B. And then, 3 will go to C. And then, 2 will go to B, so, 1 will go from A to C." (lines 165-169). At the last major move, her subgoals are pyramids: "Next, if the three at A go to C, I will be done. So first, the top two disks will be moved to B. For that, 1 goes from A to C."

At the 5 major moves in episode 4, the subject uses disk-based subgoal on the first one (lines 164-169), rule 6 on the second one (lines 173-174), and pyramid subgoal on the rest. At the first appearance of the pyramid subgoal strategy, the subject says:

- | | | |
|-----------|------|--|
| [5,4,123] | 178. | Next, 5 has to go to C, so... |
| | 179. | I only need move three blocking disks to... B. |
| | 180. | So, first...1 will go from C to B. |

Although the pyramid strategy appears immediately after a pause at the end of line 178, there is no reason for an impasse to occur here because disk subgoal will handle the situation perfectly. Indeed, exactly the same puzzle state was handled by disk subgoal in the preceding episode (lines 128-132). This leads me to concur with Anzai and Simon (1979), who conclude:

The new strategy appears to have been arrived at less systematically--certainly with less awareness--than the goal recursion strategy. Perhaps it would not be misleading to call the learning process here perceptual. The subject learned to view the goals in a new way--requiring not the transfer of a succession of disks but the transfer of a pyramid of disks. Unfortunately, in this task environment as in others that have

been studied, the verbal protocols give us only the slightest hints of the perceptual processes and the perceptual learning that may be going on. [pp. 127-128]

By the way, the analyses of Anzai and Simon agree with the analyses offered above, except that they are less specific. See (VanLehn, 1989) for an explicit comparison.

Although the generation of the pyramid strategy seems to take place without awareness, the subject seems to realize that it is a different strategy. On the major move just following the one where she first exhibits the pyramid strategy, the subject seems to explicitly compare its predictions with those of the old strategy:

- [_,1234,5]
- 192. 5 is already at C, so...
 - 193. I will move the remaining four from B to C...
 - 194. It's just like moving four, isn't it?
 - 195. So... I will have to move 4 from B to C...
 - 196. For that, the three that are on top have to go from B to A...
 - 197. Oh, yeah, 3 goes from B to A!
 - 198. For that, 2 has to go from B to C.
 - 199. for that, 1 has to go from B to A.
 - 200. So, 1 will go from B to A.

The subject seems to apply the pyramid rule in line 193 and 196, then notices that the disk rule yields the same results (line 197). Thereafter, she applies the disk rule, although I strongly suspect that she is checking its results against those of the pyramid rule as she goes.

In short, the triggering event and reasoning for the initial use of the pyramid rule are unknown, although probably perceptual in character. The second use of the strategy somehow triggers another learning event, at line 197, where the subject more-or-less explicitly compares disk subgoalings with pyramid subgoalings. Neither learning events corresponds to an application of standard machine learning techniques.

4. Classification of the learning events

According to the analysis above, there are four major pieces of strategic knowledge acquired during this protocol: (1) learning rule 5, which is that getting disk 4 on peg B is a prerequisite for the top level goals of the puzzle, (2) learning rule 6, which moves two-high pyramids out of the way, (3) learning the disk subgoalings strategy (rules 7 through 11), and (4) learning the pyramid subgoalings strategy (rule 12). Each will be summarized in turn.

The analysis will be summarized along two dimensions: What triggers the learning event? What types of reasoning occurs during the learning event?

The learning event for rule 5 is triggered by an impasse and the reasoning during it seems to be deductive. The impasse occurs when the existing rules recommend moving 4 to B, but that move cannot be made legally. The subject uses rule 4 to deduce that this goal will always be a prerequisite for achieving the initial top level goal (moving 5 to C), so she adds 4-to-B as a top level goal. Thus, this learning event can be classified as impasse-driven explanation-based learning.

Rule 6 also seems to be learned at an impasse. The impasse occurs because the existing rules do not uniquely determine a move. The subject seems to search forward in her imagination in order to ascertain which of the two proposed moves is better, then forms a rule that records what she has discovered. This type of reasoning does not fit the classic mold of explanation-based learning, for there is no sign of deduction from general rules. On the other hand, the reasoning is not much like similarity-based learning, for there is no induction over multiple exemplars. The reasoning seems to best fit a type of learning called *patching*, which was invented by Brown and VanLehn (1980) to explain how students acquire stable buggy strategies by encoding the results of applying a repair strategy to an impasse. In this case, patching led to the acquisition of a correct strategy rather than a buggy one. So this learning

event can be classified as impasse-driven patching.

The disk subgoal strategy is acquired over a series of five learning events. The initial learning event seems to be quite different from the subsequent ones. It occurs at line 78, while the subject is reflecting on the solution she has just made to the two-disk puzzle. She seems to explain her move to herself, deducing from general principles that it is an appropriate strategy in the given circumstances. Thus, this learning event could be classified as explanation-based learning triggered by a deliberate plan of reflecting on the solution of a simpler version of the puzzle.

It is fairly clear that the initial version of the rule is overly specific. At exactly the points where an overly specific rule would fail to apply, the subject shows signs of impasses. There are four such occasions. In all cases, she generalizes the rule just enough to get it to match the situation present at the impasse. Thus, it takes four impasses to learn a fully general rule. This conservative, gradual generalization of the rule seems to be a clear case of induction. So these four learning events can be classified as impasse-driven similarity-based learning.

The pyramid strategy seems to appear without any impasse. It seems to appear in a fully general form, which is characteristic of explanation-based learning. However, explanation-based learning is not indicated in this case since the pyramid rule is not a deductive consequence of the existing rules. As Anzai and Simon suggested, it may be that the subject's learning is just the simple substitution of the perceptually more salient feature of "pyramid" for "disk" in the old disk subgoal strategy. However, the nature of protocol data makes it difficult to tell if their suggestion is correct. This learning event's triggering and reasoning are unknown.

Oddly, on the subject's second use of the pyramid subgoal strategy, there seems to be a second learning event. At line 197, she interrupts her use of the pyramid rule and starts using the old disk subgoal strategy. This suggests that she is deliberately comparing the two strategy's execution by running the disk subgoal strategy overtly while covertly running the pyramid subgoal strategy. While this kind of reasoning is rational enough, it does not fit any of the established forms of learning in the machine learning literature.

A major feature of the protocol, which has not yet been discussed, is the subject's "experiment" of successively solving larger puzzles, presumably in order to find a general solution strategy. In a sense, this whole experiment constitutes a sophisticated learning event. For instance, she did not just solve the simpler puzzles, but seemed to pause after each in order to see if she could infer something from their solution (and she succeeded, for this was how she found the initial version of the disk subgoal strategy). As more evidence for the sophistication of her experiment, there are signs that she deliberately ignored rule 6 in order to find a more general rule. This was a fortunate choice, for rule 6, when used in combination with rules 1 through 5, suffices to solve any puzzle smaller than five disks. Had she not ignored rule 6, she may never have suffered the impasses that seem to be crucial for acquiring a general rule. To my knowledge, no machine learning program has demonstrated such sophistication in its approach to strategy acquisition.

The trigger for this extended learning event is not clear. The minor move strategy is not sufficient to determine the initial move of the five-disk puzzle, so it might be an impasse that causes her to implement the experiment. On the other hand, there is no sign of pauses or confusion prior to the initiation of the experiment (lines 70-74). Instead, the experiment seems to be triggered by curiosity, for the subject says "I wonder if I've found something new..." An interpretation for her behavior will be offered in a moment.

5. Conclusions

With respect to triggering of learning events, current machine learning techniques suffice for 7 of the 10 learning events. Six learning events are triggered by impasses, and one learning event (the acquisition of rule 7, the initial version of the disk subgoal strategy) seems to be triggered by a deliberate goal of reflecting on a solution. Impasses and deliberate retrospection are well known

techniques. For instance, Soar's learning is driven by impasses (Laird, Newell, & Rosenbloom, 1987) and Prodigy's learning is driven by deliberate examination of its search tree (Minton et al., 1987; Minton et al., 1989).

Of the three learning events whose triggering events do not correspond to standard machine learning techniques, two seem quite similar. One learning event is triggered by a desire to check a new strategy by comparing its recommendation to the recommendation made by an older strategy. To my knowledge, no machine learning program does this sort of checking, even though it would be easy to implement. Another learning event, arguably the most important one in the whole protocol, seems to be triggered by conjectural reasoning of some kind. After reflecting for a moment on her first successful solution attempt, the subject says "I wonder if I've found something new...." then begins her experiment of solving successively larger versions of the puzzle. My guess is that the subject has developed a conjecture of some kind involving the sizes of the disks. It is not at all clear from the protocol what her conjecture is, but whatever it is, it causes the subject to set up a classic Poylesque experiment of looking for a pattern involving the sizes of the disks. In this respect, both this learning event and the previous one treat the strategic knowledge base as a scientist treats a theory: a collection of solidly believed items mixed with a few new conjectures whose truth is to be tested by experiment. Learning by experimentation is being actively investigated by many machine learning researchers in the context of discovering qualitative physical theories (e.g., Falkenhainer and Rajamoney, 1988; Langley et al., 1987), mental models of complex devices (e.g., Dietterich and Buchanan, 1983) or properties of primitive problem solving operators (e.g., Carbonell and Gil, 1987). This is the first hint of its applicability to the discovery of problem solving strategies. The basic idea might also apply to planning and other design tasks.

With regards to the type of reasoning that goes on during a learning event, again the standard machine learning techniques suffice in most cases. There are two cases of explanation-based learning, four cases of similarity based learning, and one case of patching. As just discussed, two further cases involve designing and conducting little experiments. Thus, 7 of the 10 learning events correspond to off-the-shelf techniques, and two more correspond to a technique that is under active development. That leaves just one learning event to discuss. Rule 12, the rule for pyramid subgoalings, seems to be acquired by simply substituting the concept "pyramid" for the concept "disk" in rule 11. This perturbation of the rule seems to be completely unmotivated in that it does fix any kind of problem that she is currently experiencing. Indeed, it is not clear whether it makes her overall strategy more or less efficient. Perhaps she has a heuristic bias towards adoptions of new ideas, even if they seem to offer no improvement in themselves, because such ideas might provide a stepping stone to development of ideas that do offer improvements. This type of unmotivated perturbation to correctly functioning rules is reminiscent of Eurisko (Lenat & Brown, 1984) and the genetic algorithms work (DeJong, 1988).

The overall picture one gets is that the subject is deliberately constructing a theory about Tower of Hanoi strategies. When she detects a deficiency in her theory, usually in the form of an impasse, then she attempts to rectify it using deduction, experimentation, induction, or if all else fails, a repair strategy. She apparently has some "curiosity" demons preset to notice interesting events and propose an exploration of them. She seems to have set the noise threshold, so to speak, on her cognitive system in such a way that small perturbations are allowed to creep into the rules, which sometimes leads to unanticipated improvements. Clearly, there is no machine learning system on earth that includes all these styles of learning, and yet, there is nothing stopping us from building one.

References

- Anderson, J. R. (1983). *The Architecture of Cognition*. Cambridge, MA: Harvard.
- Anzai, Y. & Simon, H.A. (1979). The theory of learning by doing. *Psychological Review*, 86, 124-140.
- Brown, J. S. & VanLehn, K. (1980). Repair Theory: A generative theory of bugs in procedural skills. *Cognitive Science*, 4, 379-426.
- Carbonell, J.G. & Gil, Y. (1987). Learning by experimentation. In Langley, P. (Ed.), *Proceedings of the Fourth Workshop on Machine Learning*. Los Altos, CA: Morgan Kaufman.
- Chi, M.T.H., Bassok, M., Lewis, M., Reimann, P. & Glaser, R. (in press, 19??). Learning problem solving skills from studying examples. *Cognitive Science*, .
- DeJong, K. (1988). Learning with genetic algorithms: An overview. *Machine Learning*, 3(2/3), 121-138.
- Dietterich, T.G. & Buchanan, B.G. (1983). The role of experimentation in theory formation. In Michalski, R.S. (Ed.), *Proceedings of the International Machine Learning Workshop*. Los Altos, CA: Morgan Kaufman.
- Falkenhainer, B. & Rajamoney, S. (1988). The interdependence of theory formation, revision and experimentation. In Laird, J. (Ed.), *Proceedings of the Fifth International Conference on Machine Learning*. Los Altos, CA: Morgan Kaufman.
- Hunt, E.B., Marin, J. & Stone, P.J. (1966). *Experiments in Induction*. New York: Academic.
- Laird, J. E., Newell, A., and Rosenbloom, P. S. (1987). Soar: An architecture for general intelligence. *Artificial Intelligence*, 33, 1-64.
- Langley, P., Simon, H.A., Bradshaw, G.L., & Zytkow, J.M. (1987). *Scientific Discovery: Computational Explorations of the Creative Process*. Cambridge, MA: MIT Press.
- Lenat, D.B. & Brown, J.S. (1984). Why AM and Eurisko appear to work. *Artificial Intelligence*, 23(3), 269-294.
- Minton, S., Carbonell, J.G., Etzioni, O., Knoblock, C. & Kuokka, D.R. (1987). Acquiring effective search control rules: Explanation-based learning in the Prodigy system. In P. Langley (Ed.), *Proceedings of the Fourth International Workshop on Machine Learning*. Los Altos, CA: Morgan Kaufmann.
- Minton, S., Carbonell, J., Knoblock, C., Kuokka, D.R., Etzioni, O., Gil, Y. (1989). *Explanation-based learning: A problem-solving perspective* (Tech. Rep. CMU-CS-89-103). Carnegie-Mellon University, Dept. of Computer Science.
- Newell, A. (in prep., 19??). *Universal Theories of Cognition*. ?addr?: ?pub.
- Newell, A. & Simon, H. A. (1972). *Human Problem Solving*. Englewood Cliffs, NJ: Prentice-Hall.

- Smith, E.E. & Medin, D. (1981). *Categories and concepts*. Cambridge, MA: Harvard University Press.
- VanLehn, K. (1983). Human skill acquisition: Theory, model and psychological validation. In *Proceedings of AAAI-83*. Los Altos, CA: Morgan Kaufmann,
- VanLehn, K. (1987). Learning one subprocedure per lesson. *Artificial Intelligence*, 31(1), 1-40.
- VanLehn, K. (1989). *Learning events in the discovery of problem solving strategies* (Tech. Rep. PCG-17). Dept. of Psychology, Carnegie-Mellon University.
- VanLehn, K. (in press, 19??). A workbench for discovering task-specific theories of learning. In T.O'Shea & E.Scanlon (Ed.), *Proceedings of the Nato Workshop on Advanced Educational Technology*.
- VanLehn, K. (in press, 19??). *Mind Bugs: The origins of procedural misconceptions*. Cambridge, MA: MIT Press.
- VanLehn, K. & Ball, W. (in press, 19??). Teton: A large-grained architecture for studying learning. In VanLehn, K. (Ed.), *Architectures for Intelligence*. Hillsdale, NJ: Erlbaum.
- VanLehn, K., Ball, W. & Kowalski, B. (1988). *Non-LIFO execution of cognitive procedures* (Tech. Rep. PCG-15). Dept. of Psychology, Carnegie-Mellon University.

Notes

¹An impasse is defined to be a situation where the immediately available knowledge is not sufficient to uniquely determine the next action to be taken (Brown & VanLehn, 1980; Laird, Newell & Rosenbloom, 1986). Impasses are defined relative to a given problem solving architecture, which is Teton in this case, and the knowledge represented in it. Since I am interested in learning events that are "visible to the naked eye," so to speak, Teton's design and knowledge representation are oriented towards aligning impasses with signs of visible distress by subjects, such as long pauses or negative self-monitoring statements (e.g., "I'm lost.", "Huh?", "What do I do now?", or "Nuts!"). Such signs of distress may be taken as an operational definition of a Teton impasse in lieu of running the simulation itself.

²It is not clear why she fails to get the impasse a moment earlier when the puzzle is in the state [45,123,_]. The same reasoning applies -- the move of 4 to B is suggested by rule 4 strategy, but blocked by the disks on peg B. The move she actually makes, 4 to C, is suggested by rules 2 and 3.. Perhaps rules two and three have priority over the rule four so that when they uniquely determine a move, the rule four is not considered.

³The protocol is unclear. She may have actually executed the move.

⁴The existence of a "looking over" phase to the patching was not conjectured by Brown and VanLehn. It is an intriguing event, for it opens the issue of whether reflection on past actions is necessary for the formulation of abstract rules.

⁵In order to get an impasse at line 84, we must assume that rule 4 does not apply. Rule four moves single disks out of the way, so it will suffice for clearing peg A for a move disk 2 to peg B when the state is [123,_,_]. Perhaps this rule is being deliberately ignored, just as rule 6 is, in order to form a more general rule. An alternative interpretation of the protocol would be to assume that the pause at line 84 comes from unexplained causes, and that the generalization attributed to line 84 occurs somewhere else, such as line 81 or line 87.

⁶The reasoning would be: If the goal is to move a disk from peg S to peg T, and there are some disks in the way, then they must all be moved out of the way. Thus, they must all be moved to the peg that is not involved in the move. Thus, they must be placed in a pyramid on the peg that is not involved in the move. The bottom disk of that pyramid must be moved first to that peg. The bottom disk is the biggest disk of the ones blocking the initial goal move. Summing up: if the goal is to move a disk from a peg to another peg, then move the largest blocking disk to the peg that is not involved in the move.